



US Patent & Trademark Office

[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search: ☒ The ACM Digital Library ☐ The Guide



THE ACM DIGITAL LIBRARY


[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used

instruction issue and dependency and control and buffer and switch

Found 46,618 of 132,857

Sort results by

Display results


[Save results to a Binder](#)

[Search Tips](#)
☐ Open results in a new window

[Try an Advanced Search](#)
[Try this search in The ACM Guide](#)

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale ☐ ☐ ☐ ☐ ☐

1 [Data and memory optimization techniques for embedded systems](#)

P. R. Panda, F. Catthoor, N. D. Dutt, K. Danckaert, E. Brockmeyer, C. Kulkarni, A. Vandercappelle, P. G. Kjeldsberg

 April 2001 **ACM Transactions on Design Automation of Electronic Systems (TODAES)**, Volume 6 Issue 2
Full text available: [pdf\(339.91 KB\)](#)
 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We present a survey of the state-of-the-art techniques used in performing data and memory-related optimizations in embedded systems. The optimizations are targeted directly or indirectly at the memory subsystem, and impact one or more out of three important cost metrics: area, performance, and power dissipation of the resulting implementation. We first examine architecture-independent optimizations in the form of code transformations. We next cover a broad spectrum of optimizati ...

Keywords: DRAM, SRAM, address generation, allocation, architecture exploration, code transformation, data cache, data optimization, high-level synthesis, memory architecture customization, memory power dissipation, register file, size estimation, survey

2 [I-NET mechanism for issuing multiple instructions](#)

L. Wang, C. L. Wu

November 1988 **Proceedings of the 1988 ACM/IEEE conference on Supercomputing**Full text available: [pdf\(978.12 KB\)](#)
 Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Conventional instruction issuing methods use hardware control mechanism to issue instructions in multiple-functional-unit systems. They reach physical limitations due to the complexity of issuing logic when they intend to issue multiple instructions per cycle. A new method, I-NET, is presented in this paper to overcome this shortcoming. I-NET uses a post-compiler to detect the data dependencies among instructions. The detected data dependence is then attached to the instruction code to form ...

3 [System-level power optimization: techniques and tools](#)

Luca Benini, Giovanni de Micheli

 April 2000 **ACM Transactions on Design Automation of Electronic Systems (TODAES)**, Volume 5 Issue 2
Full text available: [pdf\(385.22 KB\)](#)
 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This tutorial surveys design methods for energy-efficient system-level design. We consider electronic systems consisting of a hardware platform and software layers. We consider the three major constituents of hardware that consume energy, namely computation,

communication, and storage units, and we review methods of reducing their energy consumption. We also study models for analyzing the energy cost of software, and methods for energy-efficient software design and compilation. This survey ...

4 An elementary processor architecture with simultaneous instruction issuing from multiple threads

Hiroaki Hirata, Kozo Kimura, Satoshi Nagamine, Yoshiyuki Mochizuki, Akio Nishimura, Yoshimori Nakase, Teiji Nishizawa

April 1992 **ACM SIGARCH Computer Architecture News , Proceedings of the 19th annual international symposium on Computer architecture**, Volume 20 Issue 2

Full text available:  [pdf\(1.03 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In this paper, we propose a multithreaded processor architecture which improves machine throughput. In our processor architecture, instructions from different threads (not a single thread) are issued simultaneously to multiple functional units, and these instructions can begin execution unless there are functional unit conflicts. This parallel execution scheme greatly improves the utilization of the functional unit. Simulation results show that by executing two and four threads in parallel ...

5 Limits on multiple instruction issue

M. D. Smith, M. Johnson, M. A. Horowitz

April 1989 **ACM SIGARCH Computer Architecture News , Proceedings of the third international conference on Architectural support for programming languages and operating systems**, Volume 17 Issue 2

Full text available:  [pdf\(1.55 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper investigates the limitations on designing a processor which can sustain an execution rate of greater than one instruction per cycle on highly-optimized, non-scientific applications. We have used trace-driven simulations to determine that these applications contain enough instruction independence to sustain an instruction rate of about two instructions per cycle. In a straightforward implementation, cost considerations argue strongly against decoding more than two instructions in ...

6 A survey of processors with explicit multithreading

Theo Ungerer, Borut Robič, Jurij Silc

March 2003 **ACM Computing Surveys (CSUR)**, Volume 35 Issue 1

Full text available:  [pdf\(920.16 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Hardware multithreading is becoming a generally applied technique in the next generation of microprocessors. Several multithreaded processors are announced by industry or already into production in the areas of high-performance microprocessors, media, and network processors. A multithreaded processor is able to pursue two or more threads of control in parallel within the processor pipeline. The contexts of two or more threads of control are often stored in separate on-chip register sets. Unused i ...

Keywords: Blocked multithreading, interleaved multithreading, simultaneous multithreading

7 Critical issues regarding HPS, a high performance microarchitecture

Y. N. Patt, S. W. Melvin, W. M. Hwu, M. C. Shebanow

December 1985 **ACM SIGMICRO Newsletter , Proceedings of the 18th annual workshop on Microprogramming**, Volume 16 Issue 4

Full text available:  [pdf\(987.20 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

HPS is a new model for a high performance microarchitecture which is targeted for implementing very dissimilar ISP architectures. It derives its performance from executing the operations within a restricted window of a program out-of-order, asynchronously, and

concurrently whenever possible. Before the model can be reduced to an effective working implementation of a particular target architecture, several issues need to be resolved. This paper discusses these issues, both in general and in ...

8 Computing curricula 2001

September 2001 **Journal on Educational Resources in Computing (JERIC)**

Full text available:  [pdf\(613.63 KB\)](#)
 [html\(2.78 KB\)](#)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

9 Energy: efficient instruction dispatch buffer design for superscalar processors

Gurhan Kucuk, Kanad Ghose, Dimitry V. Ponomarev, Peter M. Kogge

August 2001 **Proceedings of the 2001 international symposium on Low power electronics and design**

Full text available:  [pdf\(106.85 KB\)](#)


Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

Keywords: bitline segmentation, low power comparator, low power instruction scheduling, low-power superscalar datapath

10 Unconstrained speculative execution with predicated state buffering

Hideki Ando, Chikako Nakanishi, Tetsuya Hara, Masao Nakaya

May 1995 **ACM SIGARCH Computer Architecture News , Proceedings of the 22nd annual international symposium on Computer architecture**, Volume 23 Issue 2

Full text available:  [pdf\(1.50 MB\)](#)


Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Speculative execution is execution of instructions before it is known whether these instructions should be executed. Compiler-based speculative execution has the potential to achieve both a high instruction per cycle rate and high clock rate. Pure compiler-based approaches, however, have greatly limited instruction scheduling due to a limited ability to handle side effects of speculative execution. Significant performance improvement is, thus, difficult in non-numerical applications. This paper ...

11 Empirical performance evaluation of concurrency and coherency control protocols for database sharing systems

Erhard Rahm

June 1993 **ACM Transactions on Database Systems (TODS)**, Volume 18 Issue 2

Full text available:  [pdf\(3.37 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Database Sharing (DB-sharing) refers to a general approach for building a distributed high performance transaction system. The nodes of a DB-sharing system are locally coupled via a high-speed interconnect and share a common database at the disk level. This is also known as a "shared disk" approach. We compare database sharing with the database partitioning (shared nothing) approach and discuss the functional DBMS components that require new and coordinated solutions for DB-shar ...

Keywords: coherency control, concurrency control, database partitioning, database sharing, performance analysis, shared disk, shared nothing, trace-driven simulation

12 Tutorial: Compiling concurrent languages for sequential processors

Stephen A. Edwards

April 2003 **ACM Transactions on Design Automation of Electronic Systems (TODAES)**, Volume 8 Issue 2

Full text available:  [pdf\(771.65 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)


Embedded systems often include a traditional processor capable of executing sequential code, but both control and data-dominated tasks are often more naturally expressed using one of the many domain-specific concurrent specification languages. This article surveys a variety of techniques for translating these concurrent specifications into sequential code. The techniques address compiling a wide variety of languages, ranging from dataflow to Petri nets. Each uses a different method, to some degree ...

Keywords: Compilation, Esterel, Lustre, Petri nets, Verilog, code generation, communication, concurrency, dataflow, discrete-event, partial evaluation, sequential

13 A dynamic scheduling logic for exploiting multiple functional units in single chip multithreaded architectures

Prasad N. Golla, Eric C. Lin

February 1999 **Proceedings of the 1999 ACM symposium on Applied computing**

Full text available:  [pdf\(1.19 MB\)](#)


Additional Information: [full citation](#), [references](#), [index terms](#)

Keywords: Tomasulo's algorithm, computer architecture, microprocessor, multithreading, threaded architectures

14 Memory access buffering in multiprocessors

Michel Dubois, Christoph Scheurich, Faye Briggs

August 1998 **25 years of the international symposia on Computer architecture (selected papers)**

Full text available:  [pdf\(1.10 MB\)](#)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

15 Memory access buffering in multiprocessors

M. Dubois, C. Scheurich, F. Briggs

June 1986 **ACM SIGARCH Computer Architecture News , Proceedings of the 13th annual international symposium on Computer architecture**, Volume 14 Issue 2

Full text available:  [pdf\(943.66 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In highly-pipelined machines, instructions and data are prefetched and buffered in both the processor and the cache. This is done to reduce the average memory access latency and to take advantage of memory interleaving. Lock-up free caches are designed to avoid processor blocking on a cache miss. Write buffers are often included in a pipelined machine to avoid processor waiting on writes. In a shared memory multiprocessor, there are more advantages in buffering memory requests, since each m ...

16 Improving instruction supply efficiency in superscalar architectures using instruction trace buffers

Chih-Po Wen

April 1992 **Proceedings of the 1992 ACM/SIGAPP Symposium on Applied computing: technological challenges of the 1990's**


Full text available:  [pdf\(898.33 KB\)](#)

Additional Information: [full citation](#), [references](#), [index terms](#)

17 Improving I/O performance with a conditional store buffer

Lambert Schaelicke, Al Davis


November 1998 **Proceedings of the 31st annual ACM/IEEE international symposium on Microarchitecture**

Full text available:  [pdf\(2.53 MB\)](#)

Additional Information: [full citation](#), [references](#), [index terms](#)

18 The Clipper processor: instruction set architecture and implementation


W. Hollingsworth, H. Sachs, A. J. Smith

February 1989 **Communications of the ACM**, Volume 32 Issue 2Full text available:  [pdf\(4.67 MB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Intergraph's CLIPPER microprocessor is a high performance, three chip module that implements a new instruction set architecture designed for convenient programmability, broad functionality, and easy future expansion.

19 Multiple instruction issue in the NonStop cyclone processor


Robert W. Horst, Richard L. Harris, Robert L. Jardine

May 1990 **ACM SIGARCH Computer Architecture News , Proceedings of the 17th annual international symposium on Computer Architecture**, Volume 18 Issue 3Full text available:  [pdf\(1.06 MB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

This paper describes the architecture for issuing multiple instructions per clock in the NonStop Cyclone Processor. Pairs of instructions are fetched and decoded by a dual two-stage prefetch pipeline and passed to a dual six-stage pipeline for execution. Dynamic branch prediction is used to reduce branch penalties. A unique microcode routine for each pair is stored in the large duplexed control store. The microcode controls parallel data paths optimized for executing the most frequent instr ...

20 EXPLORER: a retargetable and visualization-based trace-driven simulator for superscalar processors

Trung A. Diep, John P. Shen, Mike Phillip

December 1993 **Proceedings of the 26th annual international symposium on Microarchitecture**Full text available:  [pdf\(1.56 MB\)](#)Additional Information: [full citation](#), [references](#), [citations](#)

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2004 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)